CUSTOMER SATISFACTION THROUGH OPTIMAL DISTRIBUTION OF STOCK

H. Laurie* and S Smith†

Abstract

We present a simulation study of out-of-stock events that (a) explains reported shortages at a redistribution centre, and (b) can be used to study the out-of-stock behaviour of various stock replenishment policies and retail demand models. Our model is conceptual and is illustrated for the very simple case of a single stock, a data-driven demand model and a trigger model of replenishment orders. Further work is needed to make this model applicable in detail to stock management.

1 Introduction

The stock distribution problem as presented contrasted that of urban versus rural customers. The stock distributor is Amalgamated Beverage Industries (henceforth ABI), the urban area is Richards Bay, the rural area is Mkhuze (approximately 160 km away), and the customers are the retailers that sell ABI's beverages. For ABI, Richards Bay is an important distribution centre. It features large warehousing and repacking capacities and routinely schedules deliveries by a large number of vehicles. However, manufacture takes place in Durban, and it was clear that deliveries directly to Mkhuze offered substantial potential savings on transport costs. Mkhuze is also a large and growing market. It seemed clear to ABI that some sort of delivery node should be developed at Mkhuze.

However, it was not clear whether it should be a fully fledged distribution centre as at Richards Bay or a smaller facility with simpler tasks. Their presentation to MISGSA2005 on the Monday morning focused on the issues

*Department of Mathematics and Applied Mathematics, University of Cape Town, Rondebosch, 7701, South Africa, *e-mail: henri@maths.uct.ac.za*

†Department of Mathematics and Applied Mathematics, University of Cape Town, South Africa.

involved in choosing the size and type of the Mkhuze facility and also on the importance to ABI of avoiding out-of-stock events.

During the week the problem was explored further in extensive discussions that included Kobus Fourie of ABI and Kevin Hingst of Rieger Industrial Consultants. It became clear that ABI had already started operations in Mkhuze, apparently using rented warehouse space. Contrary to both policy and practice of ABI, they experienced out-of-stock events in this operation. Since the rural setting was novel for ABI, they wondered whether perhaps their approach to stocking levels at the Mkhuze warehouse was at fault. In particular, their restocking policy was based on expected demands. They wondered whether the expected demands generated for use in this policy was occasionally too low for the rural area of Mkhuze.

We decided to focus on the out-of-stock aspect of the problem, and report on it here. Our report is in three parts. In the first part, we define the problem mathematically, and show that the observed out-of-stock events at Mkhuze should *not* be attributed to the restocking policy, but rather to the fact that recommended replenishment order levels could not be delivered, principally because of warehouse and transport constraints. In the second part, we analyse some ABI demand data in detail and do a data-driven simulation study for a simple implementation of the mathematical model. The purpose here is to show how the model can used to estimate the frequency of out-of-stock events, given a demand model and a restocking policy. In the third part, we discuss further work, which concern more sophisticated options of the model, alternative approaches to simulation and the use of data, and applications.

2 The mathematical model

An out-of-stock event occurs when the demand for local deliveries exceeds what can be supplied by the warehouse. We will use the following notation:

$s(i)$ — the amount of stock in the warehouse at the start of day i

$d(i)$ — the amount of restocking deliveries received by the warehouse during day i

$x(i)$ — the demand on the warehouse for local deliveries during day i

By implication, we are assuming that these three variables describe all movement of stock. In particular, there are no incorrect deliveries, returns or disposing of stock past its sell-by date.

On day i, the maximum demand that can be met is $s(i) + d(i)$. We interpret this to mean that $x(i) <= s(i) + d(i)$. In other words, for us $x(i)$ represents the deliveries as achieved by ABI rather than as requested by customers. This leads to a basic balance law:

$$s(i+1) = \max \begin{Bmatrix} s(i) + d(i) - x(i) \\ 0 \end{Bmatrix} .$$

In order to complete the model, we need to specify achieved consumer demand $x(i)$ and factory deliveries $d(i)$. There are many ways to do so. For us, $x(i)$ is randomly generated in a way consistent with demand data. In Section 2 we do this in a direct way and in Section 3 we consider indirect ways that may be better. It takes a bit longer to describe $d(i)$. The factory deliveries during day i take place because of a replenishment order placed l days earlier. We denote by $o(i)$ the replenishment order placed on day i, and we assume that all replenishment orders are exactly filled, so that $d(i) = o(i - l)$. But o is by no means random; it is calculated according to a restocking policy that takes into account current warehouse stock, the expected demands between the day of the order and the day of delivery, the expected replenishments over the same period and the desired level of stock at the end of the day of delivery. We will use the following notation:

l — the delay between the day a replenishment order is placed and the day it arrives

$o(i)$ — the replenishment order placed on day i

$x_{\text{pred}}(i + j)$ — the prediction on day i of the demand during the j-th day later

$s_{\text{pred}}(i + j)$ — the warehouse stock predicted on day i at the end of j days later

$s^*(i)$ — the desired level of warehouse stock at the end of day i

$X(i, k)$ — the set of predicted demands from day i to day $i + k$

$O(i, k)$ — the set of replenishment orders placed on days i to $i + k$.

Then in general,

$$o(i) = f(s(i), \quad s^*(i + l), \quad X(i, l), \quad O(i - l, i - 1)) .$$

In the case of ABI, s^* is the expected demand over the n days after day $i+l$, that is,

$$s^*(i+l) = \sum_{j=1}^{n} x_{\text{pred}}(i+l+j)\ .$$

We see that it depends on $X(i+l,n)$. This is considered further in Section 3. Combining what we have so far, we obtain

$$\begin{array}{rcl} s(i+1) & = & s(i) + o(i-l) - x(i) \\ o(i) & = & f(s(i), s^*(i+l), X(i,l), O(i-l, l-1)) \\ x(i): & & \text{random variable model of the demands;} \\ & & \text{must satisfy } x(i) <= s(i) + o(i-l) \\ f(s, s^*, X, O): & & \text{a function determined by restocking policies} \end{array}$$

The model is presented in this form to make explicit the role of predicted demands in the replenishment orders in ABI's stock distribution policy. This policy determines what f to use, and one can experiment with variations in policy by using different f. For the sake of simplicity we sacrifice some realism; in particular, further constraints are imposed by warehouse and/or transport limitations but are ignored here.

We should note that the above model achieves some of its simplicity of representation at the cost of increased mathematical complexity. This is because x depends on the realised value of s and o which in turn depend on x, so that the model is both implicit and non-linear. We consider an alternative formulation in the last section.

To sum up, the warehouse model is fully specified by the balance law, the demand model and the restocking policy. The demand model has a stochastic component, hence stock level s is a random variable. In principle, we must expect that the probability of $s(i) = 0$ cannot be ignored. This is the definition of an out-of-stock event, of course. Our model assigns a probability to out-of-stock events, and while this may be very small, it will usually not be zero.

In the next section we illustrate the model by means of a very simple restocking policy and a demand model consistent with real data.

2.1 An example

We consider the daily demands for '1.25 Litre Coke' (Product Code 1012) at the Richards Bay distribution centre from 5 March 2002 to 26 January

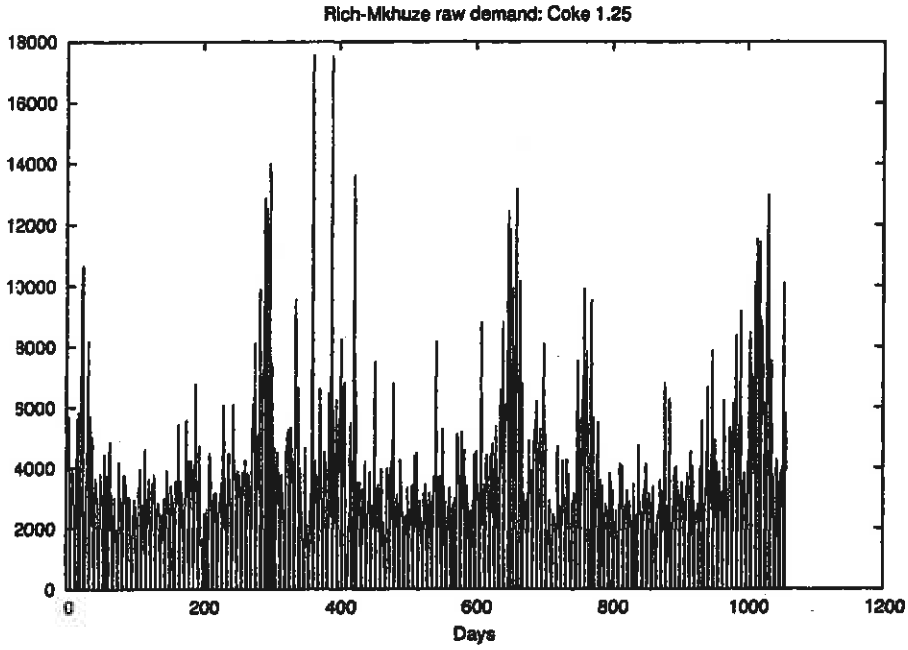

Figure 1: Raw demand data, showing total delivery to retail clients per day. The illustration represents the daily demand for '1.25 Litre Coke' (Product Code 1012) at the Richards Bay distribution centre.

2005. These form vector y. We wish to define a stochastic process that generates daily demand $x(i)$ in a way consistent with y. To do so, we need a statistical description of y.

Let us start with some exploratory analysis of the raw data, which are plotted in Figure 1. We notice occasional large values and a regular pattern of zeroes.

These zeroes are easily explained: they are an artefact concerning weekends. Deliveries to local customers over weekends are recorded as part of the deliveries on the Monday before or the Friday after. Other zeroes appear to correspond to public holidays. Hence the zeroes are not part of the demand pattern and can simply be removed from the data.

As for the large values, the histogram of the data without the zeroes shows that they are unlikely to be bad data (see Figure 2), so they were retained.

The time series and histogram of the clean data are shown in Figures 2 and 3. It is not hard to simulate a time series with statistically the same frequency histogram as y (for example by drawing randomly from y itself), but this is unsatisfactory because autocorrelation in the demands can be expected. And indeed, the partial autocorrelation function in Figure 4 for y reveals significant autocorrelation at least at lag 1, lag 2, and lag 3.

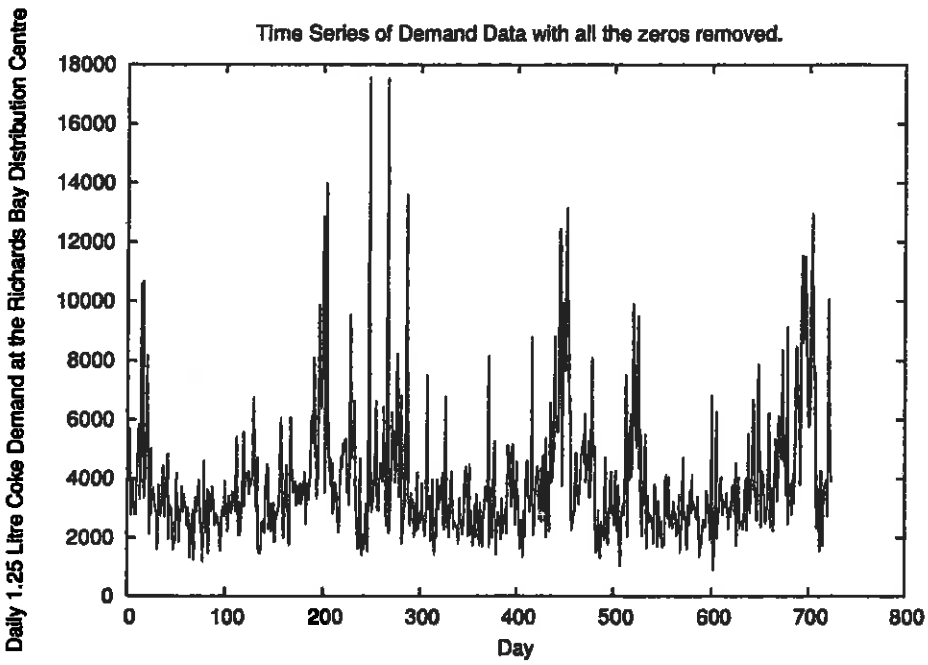

Figure 2: Time series of clean demand data: zeroes removed.

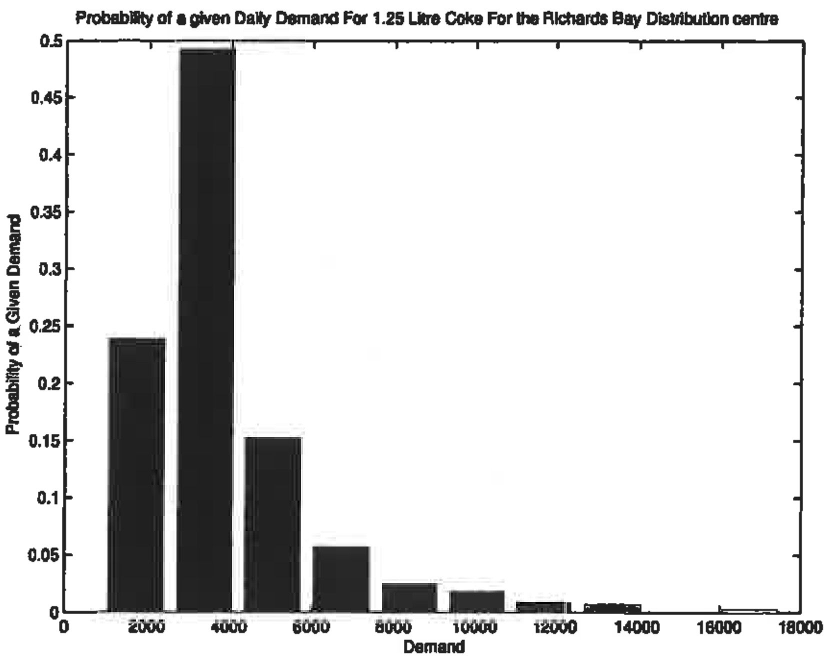

Figure 3: Histogram of clean demand data

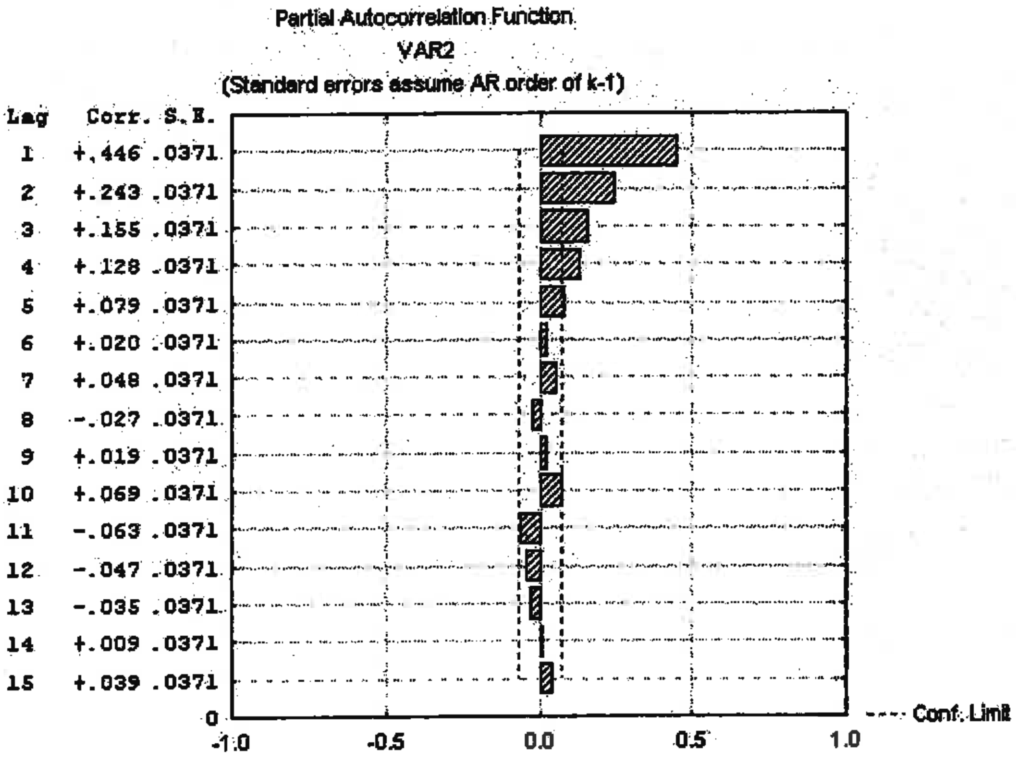

Figure 4: Partial autocorrelation function of clean demand data. Note the significant autocorrelations at lags 1, 2 and 3.

We generated random demands autocorrelated at lag 1 as follows. For all $y(i)$ in one bin of the histogram of y, we constructed the set of corresponding $y(i+1)$ values. The simulation is then simply to draw $x(i+1)$ randomly from the set corresponding to $x(i)$. We adjusted the bin width so that the sets corresponding to a each bin were of similar sizes. In other words, all bins contained more or less the same number of values. The partial autocorrelation function of a time series generated in this way appears in Figure 5 and we see that as expected it matches the partial autocorrelation function of y at lag 1, but no further. Unfortunately, this rather crude method does not extend easily to higher lags (unless a much longer training set of data is available), and in the final section we consider the future work that should be done to simulate x with the required autocorrelation also at the higher lags.

There is an additional requirement that any $x(i)$ must satisfy, namely the constraint that $x(i) <= s(i) + o(i-l)$. We deal with this as follows:

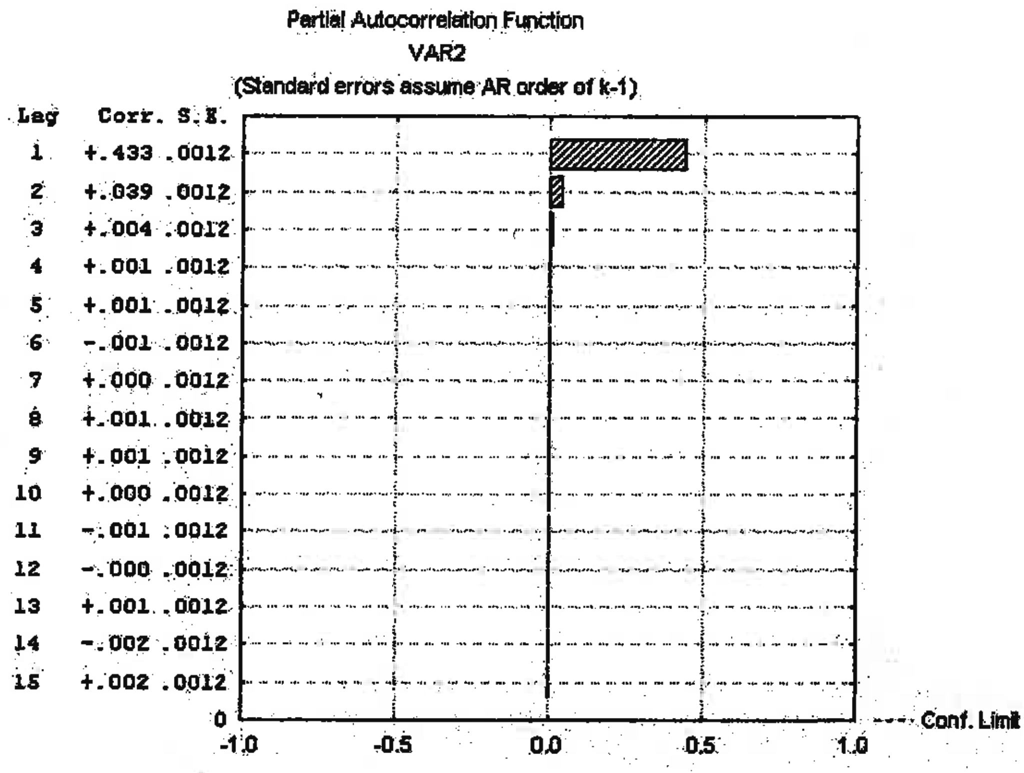

Figure 5: Partial autocorrelation function of simulated demand data.

when an $x(i)$ violating the inequality is generated, it is discarded and $x(i) = s(i)+o(i-l)$ is used instead, and an out-of-stock event is recorded. Of course, this modifies the frequency distribution of the $x(i)$. We assume that for low frequencies of out-of-stock events these departures have minimal effect on the frequency distribution, and indeed *post hoc* examination of the actual x used confirms this—for an example, see Figures 6 and 7.

We note that the model as described here differs from standard inventory control approaches [1, 2]. This is because it was designed specifically to simulate the ABI approach to stock distribution and to serve as an exploration tool. We now demonstrate a few case studies of this approach.

2.2 Test case: order double the stock shortfall.

A simple management procedure is the following: pick a target stock level s^*, and set

$$o(s) = \begin{cases} 0 & \text{if } s \geq s^* \\ 2(s^* - s) & \textit{otherwise.} \end{cases}$$

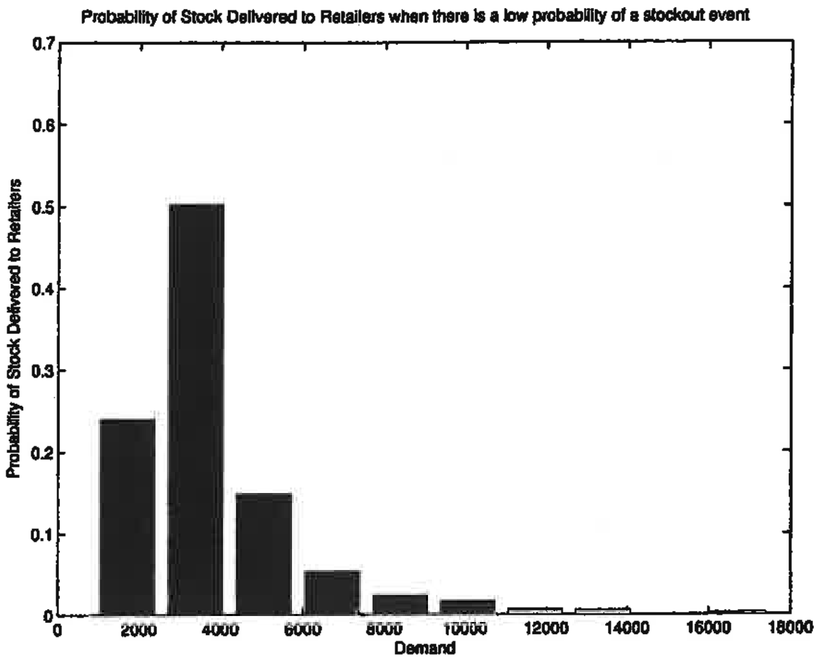

Figure 6: Histogram of typical simulated demand.

Of course, with lag $l > 1$ and autocorrelated demands, this is not a very good policy, because a large demand on day i is likely to be followed by large demands on several subsequent days, which this policy cannot anticipate.

But an improved policy would play exactly the same role as this one in any simulation based on our model, so for purposes of illustration it is adequate.

We run our simulations over 200 days. For each 200 day period, we count the number of out-of-stock events. Over many simulations, we build up a probability distribution for the number of out-of-stock events in 200 simulation days. In the limit of infinitely many simulations, this distribution depends only on l, s^* and $s(0)$. We now turn to how we handled these parameters.

The management policy as stated allows a maximum order of $2s^*$, corresponding to the order on the day after an out-of-stock event. We arbitrarily selected 80 000 as a good maximum order consistent with the data (the reader is invited to refer back to Figure 2) and hence for our simulations we used $s^* = 40\ 000$.

We then considered $s(0)$. With $l = 0$ and $s^* = 40\ 000$, the histogram of out-of-stock events shows little change for values of $s(0) > 25\ 000$ (data not shown). Since we are not interested in the effects of small $s(0)$, which affect just the start-up, we chose that level.

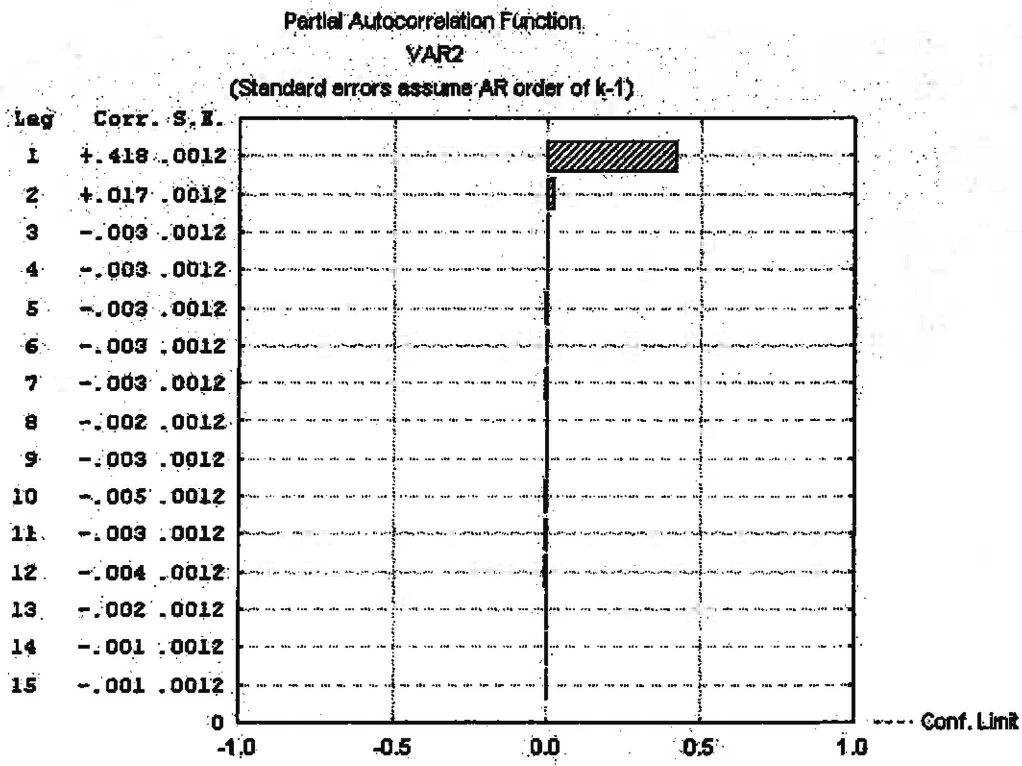

Figure 7: Typical simulated data.

Finally, we felt that the effect of varying l was instructive. Accordingly, our results concern simulations run over 200 days with $s(0) = 25\ 000$, $s^* = 40\ 000$ and $l = 1, 2$ and 3. The corresponding histograms for the proportion of out-of-stock events are shown in Figure 8. As expected, this particular management procedure is not particularly good at avoiding out-of-stock events.

3 Further work

The work above indicates that it is possible to simulate the effectiveness of a given management policy in avoiding out of stock events. However, it is merely proof of concept for this approach to out of stock events in ABI warehouses, not a full model. We address here first the further work needed to achieve that, and then make a few general comments on how such a model may be used.

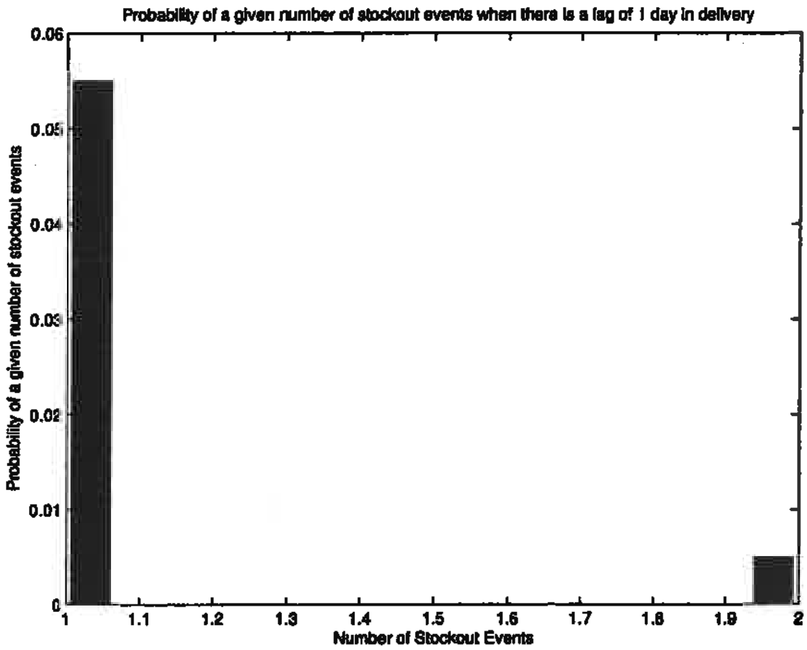
Probability of a given number of stockout events when there is a lag of 1 day in delivery
Probability of a given number of stockout events
Number of Stockout Events

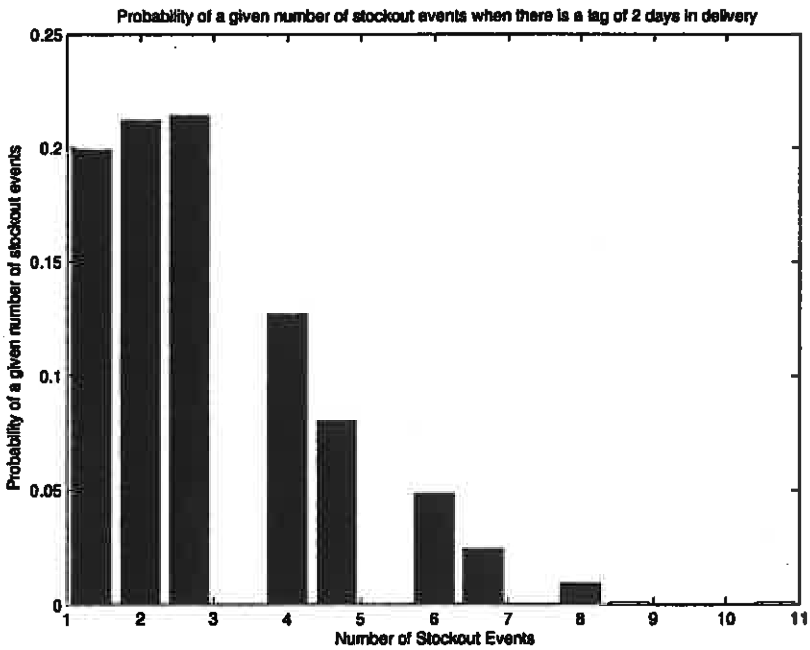
Probability of a given number of stockout events when there is a lag of 2 days in delivery
Probability of a given number of stockout events
Number of Stockout Events

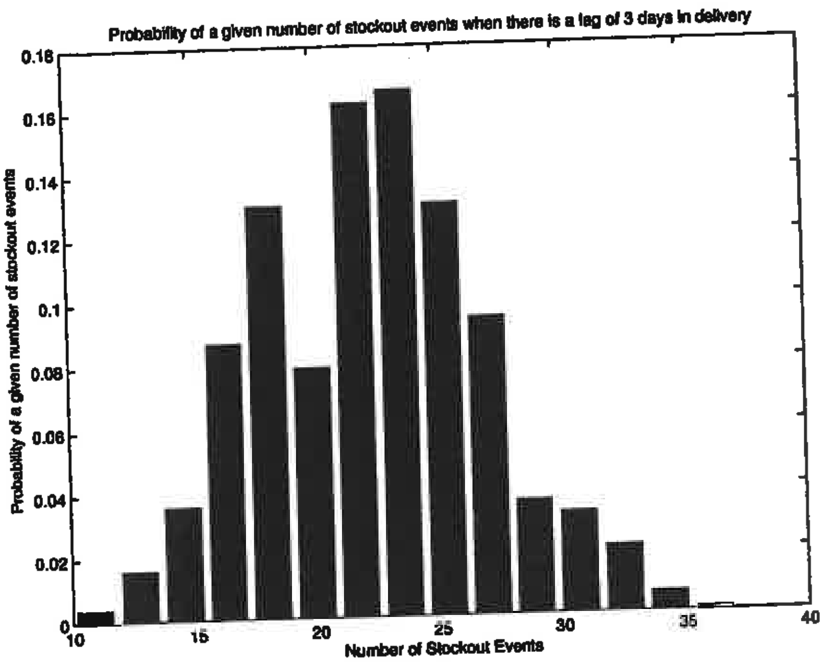

Figure 8: Histograms showing proportions of out-of-stock event numbers for various lags in the illustrative model. In all cases, $s(0) = 25\ 000$ and $s^* = 40\ 000$. The horizontal axis refers to the number of out-of-stock events in one 200 day simulation, and the vertical axis to the proportion of such simulations with that number of out-of-stock events. Zero out-of-stock events are not included, hence the vertical scale.

3.1 Matching the autocorrelation up to lag 3

Our approach to the simulation was based on the idea that the generation of a time series with a given autocorrelation signature is an unsolved problem (Colin Tredoux, personal communication). However we consider that the following approach has promise for the problem at hand:

1. generate a proposed 3-dimensional probability distribution P from which $x(i)$ can be drawn, given $x(i-1)$, $x(i-2)$ and $x(i-3)$
2. tune the parameters of P to achieve a satisfactory match with the data with respect to the frequency distribution and partial autocorrelation function of demands
3. use P to generate the demands.

3.2 More than one stock

Obviously, the work above can and should be directly extended to cover other stock-keeping units. The minimum work in doing so is writing the program code in such a way that it can accept an arbitrary input file of time series data. Then, for a given stock-keeping unit, it is just a matter of cleaning up the data and running the simulations. Cleaning the data which will be harder in some cases—we used a data set with no obvious anomalies apart from the zeroes. However, we did note that some of the time series included negative demands. Apparently these are a bookkeeping device, and do not affect the accuracy of the cumulative demand except on a very small number of days.

An interesting possibility is to use the aggregate demand also. It is likely that there is correlation between demands, so that aggregate demand could be predicted with greater accuracy than individual demands. Thus one might be able to ensure the availability of second-choice stocks to customers despite a relatively large probability of occasional out-of-stock events in individual stocks.

3.3 Greater realism and fit with ABI

The most obvious task towards a model that actually relates to the ABI out-of-stock problem is that the ABI management procedure should be used. But this should be straightforward, if (as seems to be the case) this procedure is very clearly understood. In fact, it is already implemented in computer code, so it should translate directly.

However, no management procedure covers all possible eventualities (and indeed attempting to do so is probably a bad idea), so even with the ABI procedure the model will be an idealisation. It is important to realise that a model is after all just a thought experiment, and is only valid insofar as its assumptions are valid.

This increase in realism should extend to the constraints imposed by warehouse capacity, maximum daily delivery and the fact that over-age stock is discarded.

3.4 Model validation

Of course, any attempt at a realistic model of ABI should be tested as rigourously as is feasible. In particular, a training data–test data approach would go a long way towards indicating whether the model can do more than just describe the data on which it is based. Only after testing could one say whether the model might be useful in practice.

Of course, ABI has much more data than has been considered in this study. The model we propose should apply to all of their stock-keeping units. The only requirements are a time series of consumer demands, the management procedure and the logistic constraints (warehouse and transport capacities, minimum lags between placing and arrival of replenishment orders. The model should also be tested on some of them, where in particular the autocorrelation signature may be different.

3.5 Extensions and applications

We note that if the tests on the model proposed in this section show that it produces reliable results, then the model can be extended to quite a variety of uses apart from estimating the likelihood of out of stock events.

These include

- Assessing the impact of growth in consumer demand. In particular, a realistic lag-3 autocorrelated model for consumer demands is likely to be scalable while remaining realistic. Since the constraints on deliveries and storage will be accurate, a fairly realistic picture of the response to the existing system to changes in demand should result.

- Similarly, one may keep the demands as they are currently estimated but change the management procecure and/or the constraints.

- A third kind of scenario simulation would be the choice that precipitated this study, namely a profound change in the delivery network such as a new node. After the model development outlined above, it should be relatively simple to assess the out-of-stock risk in developing a node at Mkhuze.

- Finally, realistic simulation models have often been used in management training, and the model proposed here could certainly be put to that use.

Acknowledgement

The authors acknowledge the contribution made by members of the Study Group during the meeting,especially by Isaac Takaidza of the Department of Applied Mathematics of the University of Limpopo.

References

[1] Johnson, L.A. and Montgomery, D. Operations Research in Production Planning, Scheduling and Inventory Control. Wiley (1974).

[2] Winston, W.L. Operations Research: Applications and Algorithms, (3rd edition). Duxbury Press (1994).